

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.

THIS PAGE BLANK (USPTO)



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

PCT/G399/03416

EJU

REC'D 23 NOV 1999

WIPO

PCT

3

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

98308894.9

PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

I.L.C. HATTEN-HECKMAN

DEN HAAG, DEN
THE HAGUE, 25/10/99
LA HAYE, LE

THIS PAGE BLANK (USPTO)



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.: 98308894.9
Demande n°:

Anmeldetag:
Date of filing: 30/10/98
Date de dépôt:

Anmelder:
Applicant(s):
Demandeur(s):
BRITISH TELECOMMUNICATIONS public limited company
London EC1A 7AJ
UNITED KINGDOM

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:
Data transport

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:
State:
Pays:

Tag:
Date:
Date:

Aktenzeichen:
File no.
Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:
H04L29/06

Am Anmeldetag benannte Vertragsstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE
Etats contractants désignés lors du dépôt:

Bemerkungen:
Remarks:
Remarques:

THIS PAGE BLANK (USPTO)

DATA TRANSPORT

This invention relates to transport of data over communications networks and, in particular, to transport of data encoded by layered encoding algorithms.

5 Networks based upon the Internet Protocol (IP) are being used increasingly to convey multi-media data transmissions, enabled by the use of compression algorithms to reduce data volumes to sufficiently low levels for transport over relatively low data rate network connections. However, problems remain to be overcome to achieve distribution of multi-media services, audiovisual services for
10 example, to a large number of client terminals having a variety of different capabilities for receipt of such services. In particular, some clients may have access only to limited data rate network connections enabling receipt of only low-resolution and low picture rate video. Other users may be connected to relatively high bandwidth corporate LANs and demand higher quality reception. Known
15 methods for providing different levels of service to different users include point-to-point services whereby a tailored version of a session is separately transmitted directly to each user at their network address, and "simulcast" techniques whereby a number of different data rate transmissions are broadcast and users may select and share that most suited to their needs. However, both point-to-point
20 and simulcast techniques involve significant overlap and duplication of data between transmitted data streams and are clearly inefficient in their consumption of network capacity.

 Layered encoding techniques, such as that implemented for example under the H.263 standard for video data compression, defined in "Video Coding for Low
25 Bit Rate Communications", International Telecommunication Union (ITU) - T Recommendation H.263, January 1998, enable data representing different resolutions of video to be encoded as separate layers of data frames. At the lowest layer, layer 0, a "lowest common denominator" encoding may be provided. Frames within layer 0 may provide a relatively low resolution representation of
30 original images, not necessarily all the original images. Data frames in higher layers may add increasing levels of detail to representations by lower layer frames or may encode images omitted from the lower layers altogether. Each layer of encoded data frames may be broadcast separately by a server, each layer to a different multi-casting network address. It is intended that most user equipment may be

able to receive the lowest layer 0 by accessing the appropriate multi-cast address for layer 0. Users who so choose, or who have equipment capability to receive higher layers may access one or more of the corresponding network addresses to enjoy a higher quality of audiovisual service. In this way, disparate client needs
5 may be satisfied by a single broadcast of each layer without unnecessary duplication of data.

Where multi-casting techniques are being used in relation to IP networks, a currently preferred protocol for transporting layers of encoded data frames is the User Datagram Protocol (UDP) as defined in "User Datagram Protocol", Internet
10 RFC 768, J. Postel, August 1980, published on the Internet by the Internet Engineering Task Force (IETF). However, while UDP offers a more rapid procedure for sending messages with a minimum of protocol mechanism, in comparison with the Transmission Control Protocol (TCP) for example, this is achieved at the expense of guaranteed delivery. Data may be lost, perhaps to the extent that a one
15 layer may be lost during conveyance over a network, or at least delayed with respect to other layers. Therefore, besides exercising a choice not to receive a higher layer of encoded data, there are involuntary reasons why a client apparatus may not receive all encoded data broadcast within a session. In both circumstances, problems may arise at a client apparatus in presenting received
20 data to a decoder in the correct order for decoding.

According to a first aspect of the present invention, there is provided a data streaming apparatus, having:

a data input for receiving data frames encoded by a layered encoding algorithm;

25 packetising means to insert received data frames, so encoded, into one or more predetermined packet structures, the data frames associated with each encoded layer being inserted into a different respective sequence of packets;

packet numbering means to assign a data sequence number to each packet generated by the packetising means, the data sequence number assigned to
30 a packet being indicative of the order of receipt, at the data input, of encoded data inserted within the packet; and

a network interface to transmit, in use, packets so created.

The present invention enables a sequence number to be assigned to each data packet, conveying encoded data, representative of the correct order for

subsequent presentation of the encoded data to a decoder. Such a sequence number enables packets received at a client apparatus to be correctly ordered, even when the client apparatus does not receive all the transmitted layers of packets or when individual packets are lost. This is particularly important where
5 differential encoding algorithms are used, such as that defined by the H.263 standard.

Differential encoders, such as those implementing H.263, generate layered data streams each having a highly variable data rate. The quantity of data required to encode each of a sequence of images differs according to the degree of
10 variation between successive images. In general, the order of output of encoded frames by an encoder is the order required for input to a decoder. However, if during transport from encoder to decoder, one layer is lost or delayed significantly with respect to another, or if particular packets are lost, problems arise at the receiving equipment in presenting the received data packets to a decoder in the
15 correct order for decoding. Therefore, while use of multi-layered encoding appears to solve the problem of accommodating different client needs, new problems arise in decoding multi-layered transmissions.

Preferably, a further sequence number may be assigned to each packet representing the order of transmission of the packet, under the control of a
20 selected protocol, within a sequence of packets conveying a particular layer of encoded data frames. Such a sequence number may be used to improve packet ordering efficiency through identifying whether all packets expected within a particular packet sequence have been received and that the next packet for decode must lies in another packet sequence.

25 According to a second aspect of the present invention, there is provided a client apparatus having:

a network interface;

packet receiving means to receive one or more sequences of data packets from the network interface, each data packet having a predetermined packet structure, each
30 of said one or more sequences of data packets conveying a different respective layer of encoded data frames generated by a layered encoding algorithm and each data packet having assigned thereto a data sequence number indicative of the order of output of encoded data, conveyed by the data packet, from said layered encoding algorithm;

packet ordering means to place said received data packets into a decoding order using said data sequence numbers; and

output means to output packets so ordered.

According to a third aspect of the present invention, there is provided
5 method of generating data packets to convey data frames encoded by a layered encoding algorithm for transmission over a communications network, each layer of encoded data frames being conveyed by a different respective sequence of data packets, including the steps of:

- (1) receiving an encoded data frame;
- 10 (2) inserting data from said data frame into one or more data packets generated according to a predetermined packet structure;
- (3) assigning, in respect of one of said one or more data packets, a data sequence number indicative of the order of receipt of encoded data inserted into said packet;
- 15 (4) writing said data sequence number at a predetermined position within said packet; and
- (5) performing steps (3) and (4) in respect of each of said one or more data packets generated at step (2).

According to a fourth aspect of the present invention, there is provided
20 method of ordering data packets received within one or more separately accessible sequences of data packets generated according to the method of Claim 5, including the steps of:

- (1) receiving one or more data packets on one or more of said one or more separately accessible sequences of data packets;
- 25 (2) selecting, from those data packets received at step (1), that data packet having the smallest assigned data sequence number amongst non-selected data packets;
- (3) outputting said selected data packet;
- (4) repeating steps (1) to (3).

30 The invention will now be described in more detail, by way of example only, with reference to the accompanying drawings of which:

Figure 1 shows a video streaming apparatus according to preferred embodiments of the present invention;

30-10-1998

Figure 2 shows a client apparatus arranged to receive signals transmitted by the apparatus of Figure 1;

Figure 3 shows a typical hierarchy of layered data frames, encoded from a small sequence of video images, for transmission by the apparatus of Figure 1;

5 Figure 4 shows the result of applying a packet numbering algorithm, according to preferred embodiments of the present invention, to packets generated to convey encoded data frames shown in Figure 3;

Figure 5 shows the structure of a packet header according to the Real-time Transport Protocol (RTP) as used in a preferred embodiment of the present
10 invention; and

Figure 6 is a flow diagram showing steps in the operation of a preferred client apparatus, relating in particular to the ordering of packets broadcast by the apparatus of Figure 1.

Preferred embodiments of the present invention will now be described
15 in the particular context of a video streaming apparatus, although the present invention may be applied to other forms of data broadcast and receiving apparatus, not necessarily in a client-server arrangement, involved in a single or multi-media session with data encoded by layered encoding techniques.

Referring to Figure 1, a video streaming apparatus 100 is shown,
20 according to preferred embodiments of the present invention, for use in broadcasting multi-layer encoded audiovisual data from an encoded audio/video data source 105 to client systems over a communications network 110. The audio/video data source 105 may for example be a database of encoded video data for use in a "video-on-demand" system, or it may be a multi-layer encoded real-
25 time audiovisual data stream to be transmitted as a live broadcast. The video streaming apparatus 100 accepts layers of encoded data from the source 105 at an input 115 before passing them to a packetiser 120. The packetiser 120 may implement a known algorithm for separately incorporating data from each layer of encoded data into a different respective stream of packets according to one or
30 more predetermined packet structures. For example, one or more layers may be incorporated into packets having a structure defined for use with the Real-Time Transport Protocol (RTP), described by Internet Request for Comment (RFC) 1889, January 1996 - "RTP: A Transport Protocol for Real-Time Applications" by H.Schulzrinne, S.Casner, R.Frederick and V.Jacobson, and published on the

Internet by the Internet Engineering Task Force (IETF). Once arranged according to their respective predetermined packet structure, the layers of packets are passed to a packet numbering module 125 to be numbered by a packet numbering method according to preferred embodiments of the present invention. The numbered
5 packets are then passed to instances of a session handler 130, one instance of session handler per layer of packets. Each instance of session handler 130 may implement an appropriate protocol to control transfer of the respective layer of packets over the communications network 110, via a network interface 135, to one or more predetermined network addresses, multi-cast addresses for example.
10 Protocols operating at lower levels in a protocol hierarchy of may be implemented by the network interface 135 as appropriate to the communications network 110. For example, at the level below RTP, the User Datagram Protocol (UDP) referenced above may be implemented by the network interface 135 to operate in conjunction with the Internet Protocol (IP).

15 Preferably, for simplicity, all layers of encoded data may be broadcast under the control, at the session level at least, of respective instances of the same protocol using the same packet structure. However, the scope of the present invention is intended to encompass those situations in which more than one type of protocol is employed to broadcast layers of encoded data received at the input
20 115.

Referring to Figure 2, a typical client apparatus 200 is shown for use in receiving, over the communications network 110, audio and/or video broadcasts by one or more sources having features in common with the video streaming apparatus 100 of Figure 1. The client apparatus 200 may create instances of a
25 session handler 210, each instance of session handler 210 "listening" for data received at a particular network address, one instance corresponding to each layer of packets received over the network via a network interface 205. The received layers of packets pass from their respective session handlers 210 into a source demultiplexer 215. In the event that multiple video streamers or other types of
30 source are broadcasting on the same session, each source may preferably be separately identifiable by the source demultiplexer 215 using information inserted into packet headers by the respective source streamer. For each distinct source identified, the source demultiplexer may create one instance of a blender 220, collating all packets received via the session handlers 210 carrying the same

source identifier, and passing the collated packets to the blender 220. The blender 220 may implement an algorithm, according to preferred embodiments of the present invention, for ordering packets received from the particular source using packet numbering information inserted by the packet numbering module 125 of the
5 respective source, video streamer 100 for example. Having established the correct packet order, taking account of any missing or inaccessible layers and packets, the blender 220 may pass the ordered packets to a depacketiser 225 to recover the layers of encoded data from the respective predetermined packet structure used by the particular streaming source, by packetiser 120 in the case of a video streamer
10 100. The depacketiser 225 passes the recovered encoded data, now in the correct sequence for decoding, to an output 230. The ordered data output from the client apparatus 200 may be taken by an appropriate decoder and, following decoding, reproduced at a display and/or audio output apparatus as appropriate.

Referring to Figure 3, a typical hierarchy of layered data frames is shown,
15 encoded from a small sequence of video images, as might be presented to the input 115 of a video streamer 100. The encoded frames 340 are shown arranged as three layers, 300, 305 and 310 corresponding to a lowest layer, a middle layer and a top layer respectively. Further layers may be generated according to the particular encoding algorithm implemented by the source 105. Each encoded frame
20 340 of Figure 3 is shown with a number in the range 1 to 10, indicating the order of output by the encoded data source 105 and hence the required order for subsequent presentation of the frames to a decoder. The frames 340 in Figure 3 are shown grouped within five columns, each column of frames being encoded to represent a respective original image 315-335. For example, original image "A"
25 315, is shown encoded as a frame number "1" in the lowest layer 300, a frame "2" in middle layer 305 and a frame "3" in the top layer 310. Original image "B" 320 is represented only in the top layer 310 by a frame generated with number "4". The original image data 315-335 would not normally be presented to the input 115 of a video streaming apparatus 100.

30 The sequence of encoded frames 1-10 of Figure 3 may, for example, be generated according to a video encoding algorithm such as H.263, referenced above. If the H.263 encoding technique is used to encode the images 315-335 of Figure 3, each frame 340 in the lowest layer 300 may represent a low-resolution version of the respective original image and may be encoded using the basic H.263

algorithm at QCIF resolution as described in Section 4.1 of the referenced specification. Frames in layers 305 and 310 represent increasingly detailed enhancements to the low-resolution image represented by the respective frame in layer 300. Under H.263, the middle and top layers may be encoded according to
5 the definition in Annex O, "Temporal, SNR and Spatial Scalability Mode", of the above-referenced H.263 specification.

Not all original images may be represented in the lower layers 300, 305. In the particular sequence shown in Figure 3, only every fourth original image is represented in the lowest layer 300 and every second original image in the middle
10 layer 305. Thus, a client apparatus able or choosing to decode only the lowest layer of frames will display a representation of the original sequence of images having a relatively low resolution and a relatively low image rate as compared with client apparatus able or choosing to decode both the lowest and middle layers. Apparatus able to receive and decode all three broadcast layers 300-310 will be
15 able to display all the original images (315-335) at the highest resolution available. It is intended that a lower data rate network connection may be used to receive data frames at the lowest layer, making that layer accessible to most client equipment.

Referring to Figure 4, a diagram is provided to show a typical breakdown
20 of those encoded frames 340, representing the first three original images 315, 320 and 325 of Figure 3, across corresponding layered sequences of packets 400 by the packetiser 120. Figure 4 also shows the result of applying a packet numbering scheme to those packets as may be implemented by the packet numbering module
25 125 according to preferred embodiments of the present invention. A typical packetiser 120 may operate to packetise each layer of encoded frames separately, generating, as in the present example, three separate streams of packets, one stream for each layer. As discussed in relation to Figure 1 above, the packetiser 120 may be arranged to implement one or more packet structures appropriate to the particular protocol chosen at the session level to control the conveyance of
30 each encoded layer of data. Preferably, each layer of encoded data may be conveyed over a network using a different respective instance of the Real-time Transport Protocol (RTP) referenced above. The packetiser 120 would, in that case, split the data within a layer of encoded frames 340 across the payload portions of a sequence of RTP packets, according to the RTP packet structure

definition. Conveniently, if packetising data encoded using the H.263 algorithm referenced above, a specific definition of an H.263 payload header is available for inclusion in RTP packets, as defined in "RTP Payload Format for H.263 Video Streams", Internet RFC 2190, September 1997, published on the Internet by the IETF. Alternative and equally satisfactory session-level protocols may be selected for implementation by the packetiser 120, employing their own respective packet structures to convey the encoded layers 300-310 of data frames 340.

Referring to Figure 4, as indicated above, each of the packets 400 is shown labelled with sequence numbers applied by packet numbering module 125. A preferred method of numbering involves the assignment of two sequence numbers to each packet 400. The number shown in the upper half of each packet 400 of Figure 4 may be referred to as a "layer sequence number" LSEQ, while the number shown in the lower half of each packet may be referred to as a "cross-layer sequence number" XSEQ. The sequence of LSEQ numbers indicates the order of transmission of packets within one specific layer. The XSEQ numbers are intended to represent the correct overall order for presentation of encoded data conveyed by those packets to a decoder 225 at a client apparatus 200. The XSEQ sequence reflects, in particular, the order that encoded data frames emerged from the source 105.

Protocols such as RTP provide a facility to assign sequence numbers to packets within a particular RTP packet stream. In this case, each layer of encoded data may be broadcast as a separate RTP packet stream under the control of different RTP session. Hence, within one layer, the respective (RTP) session handler 130 may automatically assign a layer sequence number LSEQ to each packet before transmission and write the sequence number at a predetermined position with the packet. Other types of protocol may not provide such a facility for assignment of layer sequence numbers. Hence the packet numbering module 125 may implement both layer sequence number assignment and cross-layer sequence number assignment if required.

With different layers being typically broadcast under the control of separate protocol sessions, as with RTP, there is no overall mechanism for assigning XSEQ numbers across layers. In order to assign a sequence of XSEQ numbers in particular, the packet numbering module 125 may be provided at a point immediately following the packetiser 120 and immediately before the

individual packet streams go to their respective session handlers 130 for broadcast. If required, the packet numbering module 125 may retain access to information on the order of receipt of encoded data frames at the input 115 in order to correctly assign XSEQ numbers to packets emerging from the packetiser 120. It is particularly important, when encoding data using a differential encoding algorithm such as that defined by H.263, to subsequently decode those data in the correct sequence. Assignment of an XSEQ number by the packet numbering module 125 provides a particularly convenient method of recording the correct data sequence at the packet level. Data loss or reordering of data typically occurs at the packet level. As will be discussed in the following, recording of a layer sequence number LSEQ and, in particular, a cross-layer sequence number XSEQ enables a client apparatus 200, according to preferred embodiments of the present invention, to re-order packets received out of sequence and to take account of missing packets and missing or inaccessible encoded data layers.

Referring to Figure 5, the packet header structure defined for use under RTP is shown. The RTP packet structure may be used by preferred embodiments of the present invention to record packet sequence numbers. Figure 5a shows the RTP header structure, including an optional RTP Header Extension, while Figure 5b shows the structure of the header extension itself, all details of which are described by the above-referenced RTP definition document. The RTP packet header of Figure 5a includes a Sequence Number field occupying the third and fourth octets. This field is used within the RTP protocol to record the transmission order of packets within the particular packet stream and may therefore perform the role of the layer sequence number LSEQ.

To accommodate the cross-layer sequence number XSEQ, the packet numbering module 125 may preferably use the optional RTP header extension, shown in Figure 5a at a position immediately following the "Contributing Source (CSRC) Identifiers". With this intention, the packetiser 120 may set the extension bit "X" - bit 4 of the RTP header - and include one RTP Packet Header extension, having the structure shown in Figure 5b, within each generated packet. Within each packet, the packetiser 120 may record a unique profile-specific identifier within the "profile" field of the header extension and may set the "length" field to 1, including one 32 bit "header extension" field. Such an extension field length should be adequate for use in recording XSEQ numbers generated within a typical

multi-media session. The packet numbering module 125 may then write an appropriate XSEQ number into the extension field of each packet received from the packetiser 120.

While the RTP packet structure includes fields suitable for recording assigned sequence numbers, other protocols and packet structures may not provide predetermined positions within their packets to carry sequence number information. If necessary, one or more further packet data streams may need to be generated by the packetiser 120, to be transmitted approximately in synchronisation with other packet streams, to convey sequence numbering information assigned by the packet numbering module 125 and linked, for example by a packet identifier, to packets carrying encoded data. On receipt of the "sequence number packet stream", a client apparatus may extract and use the sequence numbering information in much the same way as described below.

As discussed above in relation to the identification of a transmitting source by the source demultiplexer 215 of a client apparatus 200, for example where multiple video streamers 100 are transmitting RTP packets over the communications network 110, the "SSRC" field in the RTP header of Figure 5a may be used by an RTP session handler 130 to record within each RTP packet an identifier for the particular video streamer 100 generating the packet. The source demultiplexer 215 of a client apparatus 200 may then read the SSRC field in received packets to distinguish between packets from one video streamer and another.

Referring to Figure 6, a flow diagram is provided to show a sequence of steps in operation of an instance of a blender 220 relating to the ordering of packets, received from a particular streamer 100, numbered by a packet numbering module 125, according to preferred embodiments of the present invention. Preferably, a variable "TOP_LAYER" may be set at a predetermined value in a particular client apparatus 200, to record the highest numbered layer that the particular apparatus is set to receive and decode, either by choice or as limited by equipment capability or network connection bandwidth. The TOP_LAYER value may be set within the range 0 to n, where n is the highest numbered layer transmitted by data streaming sources accessible over the network 110.

Referring to Figure 6, processing by the blender 220 may be seen to begin at STEP 600. At STEP 602, a pre-processing step is performed on packets already

received to place them into layer sequence number (LSEQ) order within their respective layer. Ordering of packets by LSEQ number may be implemented by a known and simple ordering algorithm and, as such, further detail of STEP 602 will not be discussed in this specification.

5 At STEP 605, the blender 220 reads the first received packet on layer 0 (PKT[0]) and uses the layer sequence number (LSEQ) and cross-layer sequence number (XSEQ) contained in that packet to initialise counters LPROG[] and XPROG respectively for use in determining the next expected number in each of the packet numbering sequences. At STEP 610, variables are initialised ready for processing
10 packets from the currently selected layer, layer 0 initially. At STEP 615, an attempt is made to read the packet having the lowest layer sequence number (LSEQ) from the layer (L) currently being processed (initially layer L=0). Packets already received in time for operation of STEP 602 will have been ordered by layer sequence number so that, among those already received, the next packet read
15 from the layer L may be assumed to have the lowest LSEQ number. If, at STEP 620, a packet is available in layer L, then, at STEP 625, the cross-layer sequence number (XSEQ) in that packet is compared with the next expected cross-layer sequence number. If, at STEP 625, the current packet is the next in the cross-layer sequence then, at STEP 660, the current packet sequence numbers are used to set
20 the XPROG and LPROG[L] counter variables before, at STEP 665, that packet is forwarded to the decoder 225.

If, at STEP 620, no packet is available on the current layer, then at STEP 675 a flag is set to indicate that packets are unavailable on a layer and processing proceeds to STEP 640 to enable higher layers to be accessed in search of packets.

25 If, at STEP 625, the current packet is not the next in the cross-layer sequence, then either the next expected packet is missing from within the current layer or it lies in another layer. A following sequence of steps attempts to establish whether the next expected packet for decoding is currently missing – possibly delayed - within the current layer, or whether it may be found in another accessible
30 layer. Therefore, at STEP 630, the blender 220 first checks whether the current packet is the next expected packet within the current layer (L). If not, then at STEP 670 a flag is set to indicate that the current packet is out of sequence in its layer before processing continues to STEP 635. If, at STEP 630, the current packet was the next expected packet in its layer, then the packet having the next

expected XSEQ number must lie in another layer. However, in case the packet is soon found not to lie in a layer accessible by the client apparatus 200, or to be otherwise lost (from STEP 670), then at STEP 635 a variable recording the smallest recently encountered XSEQ number is updated along with the layer in which that respective packet was found. This will be the point of continuation in processing if the packet with the next expected XSEQ number is not located in any of the accessible layers. But first, any other accessible layers are checked.

The layer number is incremented at STEP 640. If the new current layer is at or below the highest layer accessible to the client apparatus at STEP 645, then processing returns to STEP 615 and the next expected packet is sought within that layer as described above. However, if at STEP 645, the new current layer is inaccessible then, at STEP 650, a check is made to determine whether packets are currently unavailable on any layer. If, at STEP 650, one or more layers have no received packets available, then processing restarts at STEP 610, looking again for the next expected packet, beginning at layer 0. If, at STEP 650, all layers have at least one packet available, then at STEP 652 it is determined whether or not the current packet is the next expected packet within its layer. If, at STEP 652 the current packet is correctly ordered within its layer, then the next expected packet in the XSEQ order must lie in a higher layer, one that is not accessible to the particular client apparatus 200. Therefore, the best that can be achieved is to select the packet with lowest XSEQ number on any layer. Therefore, at STEP 655, the layer having the packet with the lowest XSEQ number is selected as the new current layer, and that selected packet is treated as the next available packet for decoding. The XPROG and LPROG[L] sequence number counters are reset to the current packet values at STEP 660 and the selected packet is sent for decoding at STEP 665.

If, at STEP 652, the current packet was out of sequence within its layer, then at STEP 680, the blender 220 may optionally elect to wait for the next one or two packets to arrive on that layer for example, in case the next expected packet within the layer arrives (in which case processing would restart at STEP 610) or to continue without further delay to STEP 685 and select the layer having the packet with the lowest XSEQ number as the new current layer, and select that packet as the next available packet for decoding, proceeding to STEP 660.

After sending a packet for decoding at STEP 665, processing returns to step 610, resetting variables related to out-of-sequence processing and resetting the layer number L to 0, before continuing as described above.

It will be clear that more sophisticated processing steps may be included
5 to implement different strategies in the event that, at STEP 652, received packets within a layer are out of sequence within that layer. If the nature of the communications network 110, or the protocols selected to transfer packets across it, are such that individual packets may be delayed within a layer, then it may be beneficial to implement more sophisticated waiting algorithms if there is a
10 possibility that the expected packet may arrive later. Such a strategy is suggested in STEP 680 of Figure 6 without going into detail. Alternatively, with pure audio data for example, the effect of a lost packet may be partially overcome by inserting a duplicate of the immediately preceding packet, rather than leave a gap or risk further delay. An equivalent strategy may be available with encoded video
15 data, if manageable under the selected encoding/decoding algorithm.

Preferably, a more sophisticated algorithm may be implemented to merge the ordering of received data packets within a layer with processing steps indicated within Figure 6 from STEP 605 onwards, rather than perform pre-processing to order received data packets by LSEQ number within each layer.

CLAIMS

1. A data streaming apparatus, having:
a data input for receiving data frames encoded by a layered encoding
5 algorithm;
packetising means to insert received data frames, so encoded, into one or more predetermined packet structures, the data frames associated with each encoded layer being inserted into a different respective sequence of packets;
packet numbering means to assign a data sequence number to each
10 packet generated by the packetising means, the data sequence number assigned to a packet being indicative of the order of receipt, at the data input, of encoded data inserted within the packet; and
a network interface to transmit, in use, packets so created.
- 15 2 A data streaming apparatus according to Claim 1, wherein the packet numbering means are arranged to assign a further sequence number to each packet generated by the packetising means, said further sequence number being indicative of the position of the packet within the respective sequence of packets.
- 20 3. A data streaming apparatus according to Claim 1, wherein the packetising means are arranged to generate one or more further sequences of packets for use in conveying data sequence numbers assigned by the packet numbering means.
4. A client apparatus having:
25 a network interface;
packet receiving means to receive one or more sequences of data packets from the network interface, each data packet having a predetermined packet structure, each of said one or more sequences of data packets conveying a different respective layer of encoded data frames generated by a layered encoding algorithm and each
30 data packet having assigned thereto a data sequence number indicative of the order of output of encoded data, conveyed by the data packet, from said layered encoding algorithm;
packet ordering means to place said received data packets into a decoding order using said data sequence numbers; and

output means to output packets so ordered.

5. A method of generating data packets to convey data frames encoded by a layered encoding algorithm for transmission over a communications network, each layer of encoded data frames being conveyed by a different respective sequence of data packets, including the steps of:

- (1) receiving an encoded data frame;
- (2) inserting data from said data frame into one or more data packets generated according to a predetermined packet structure;
- 10 (3) assigning, in respect of one of said one or more data packets, a data sequence number indicative of the order of receipt of encoded data inserted into said packet;
- (4) writing said data sequence number at a predetermined position within said packet; and
- 15 (5) performing steps (3) and (4) in respect of each of said one or more data packets generated at step (2).

6. A method of generating data packets according to Claim 5, wherein step (3) includes assigning a further sequence number to said one of said one or more data packets indicative of the order of transmission of said data packet within a respective sequence of packets, and wherein step (4) includes writing said further sequence number at a further predetermined position within said data packet.

7. A method of ordering data packets received within one or more separately accessible sequences of data packets generated according to the method of Claim 5, including the steps of:

- (1) receiving one or more data packets on one or more of said one or more separately accessible sequences of data packets;
- (2) selecting, from those data packets received at step (1), that data packet having the smallest assigned data sequence number amongst non-selected data packets;
- (3) outputting said selected data packet;
- (4) repeating steps (1) to (3).

30-10-1998

EP98308894.9

CLMS

8. A method of ordering data packets within one or more separately accessible sequences of data packets received over a communications network, each sequence of data packets conveying data frames relating to a different layer of encoded data frames output by a layered encoding algorithm, each data packet
- 5 having assigned thereto a data sequence number indicative of the order of output of encoded data, conveyed by said data packet, from said encoding algorithm, and a further sequence number indicative of the position of said data packet within the respective sequence of data packets, the method comprising selecting data
- 10 packets in order of receipt within a first of said accessible sequences of data packets, outputting selected packets from said first sequence in order of assigned data sequence number and, upon selecting a packet from said first sequence having an out-of-sequence data sequence number, using the further sequence
- number assigned to said selected packet to determine whether the next expected packet, according to data sequence number, is associated with other than said first
- 15 sequence of data packets.

THIS PAGE BLANK (USPTO)



ABSTRACT
DATA TRANSPORT

A data streaming apparatus (100) is provided for broadcasting data,
5 encoded by a layered encoding algorithm, each layer of encoded data frames being
conveyed within separate respective streams of data packets of a predetermined
packet structure, including a packet numbering facility (125) to assign to each data
packet a data sequence number indicating the correct order for subsequent
decoding of encoded data conveyed by that packet. Corresponding client
10 apparatus (200) is also provided to receive data packets having such data
sequence numbers assigned to them, including a packet ordering facility (220) to
place out-of-sequence packets into the correct order for output to a decoder.

Figure 1

15

THIS PAGE BLANK (USPTO)

**Figure 1****Figure 2**




Figure 3




Figure 4

3/5

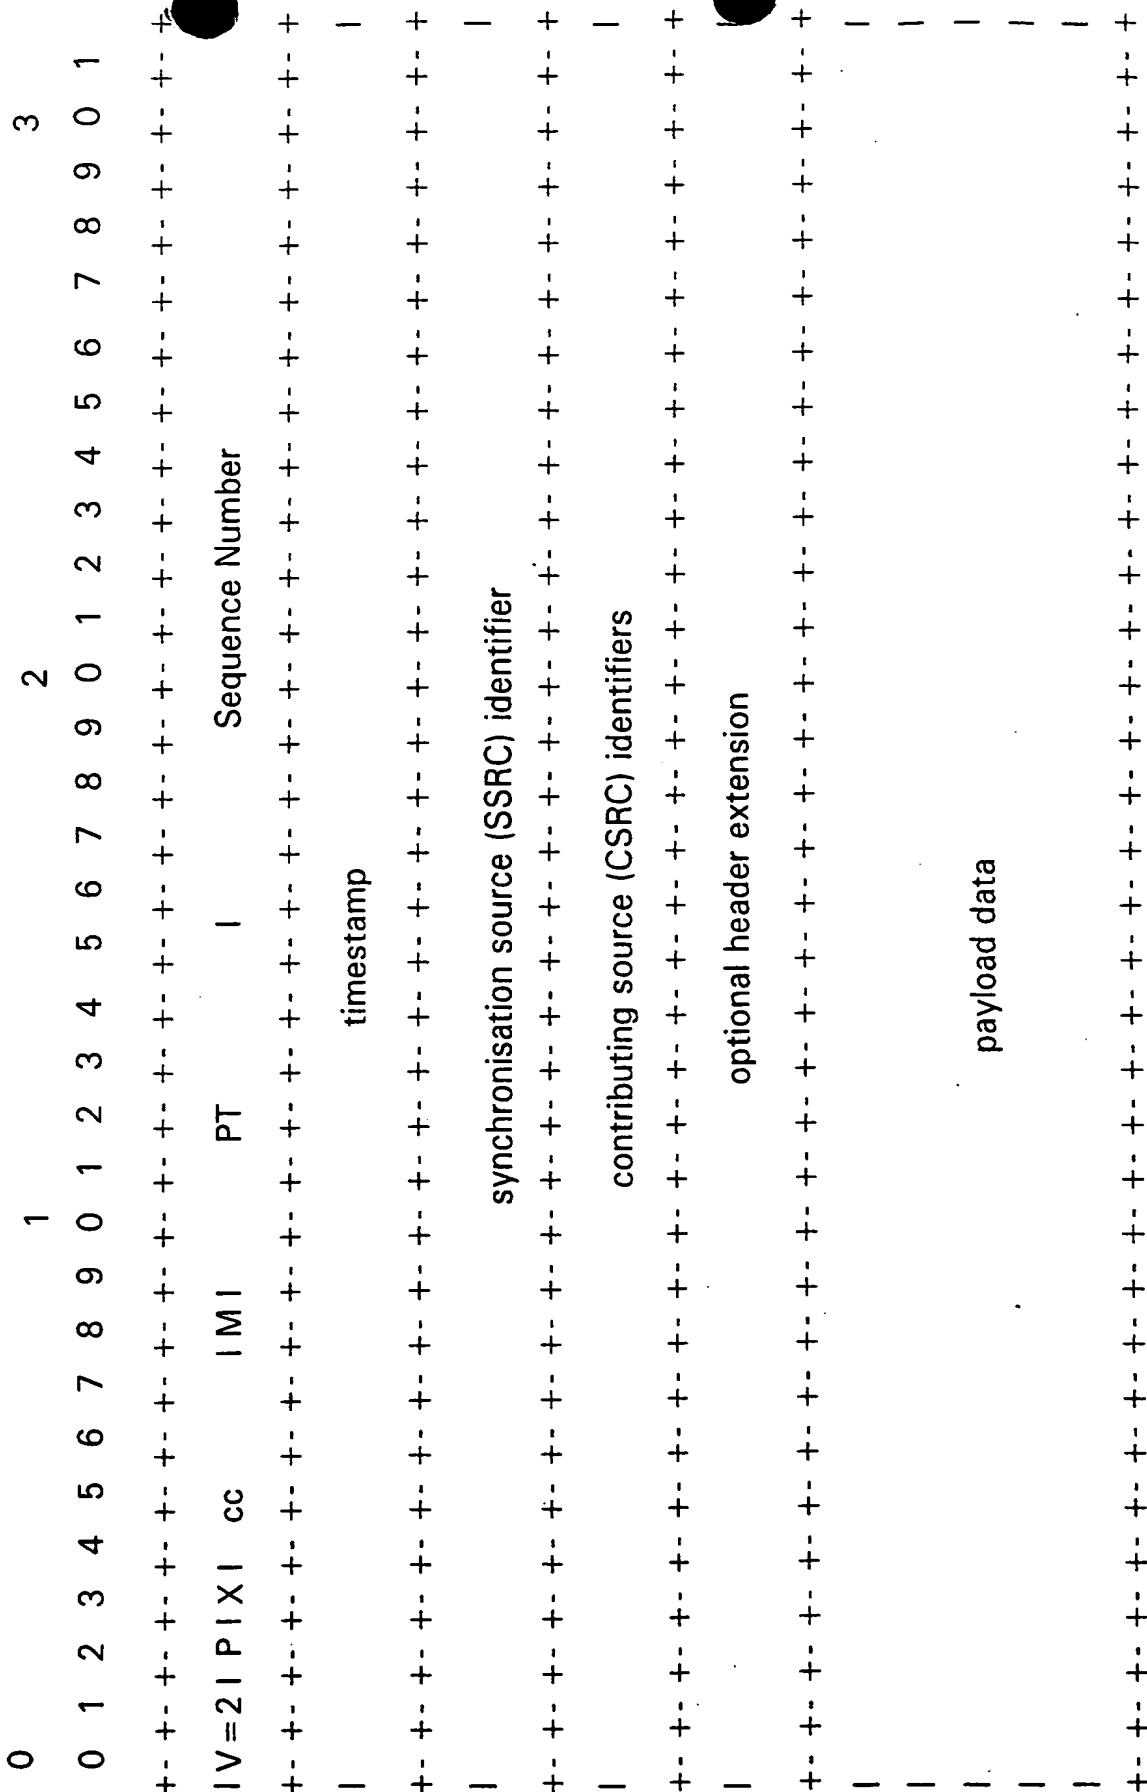


FIGURE 5a

4/5




FIGURE 5b

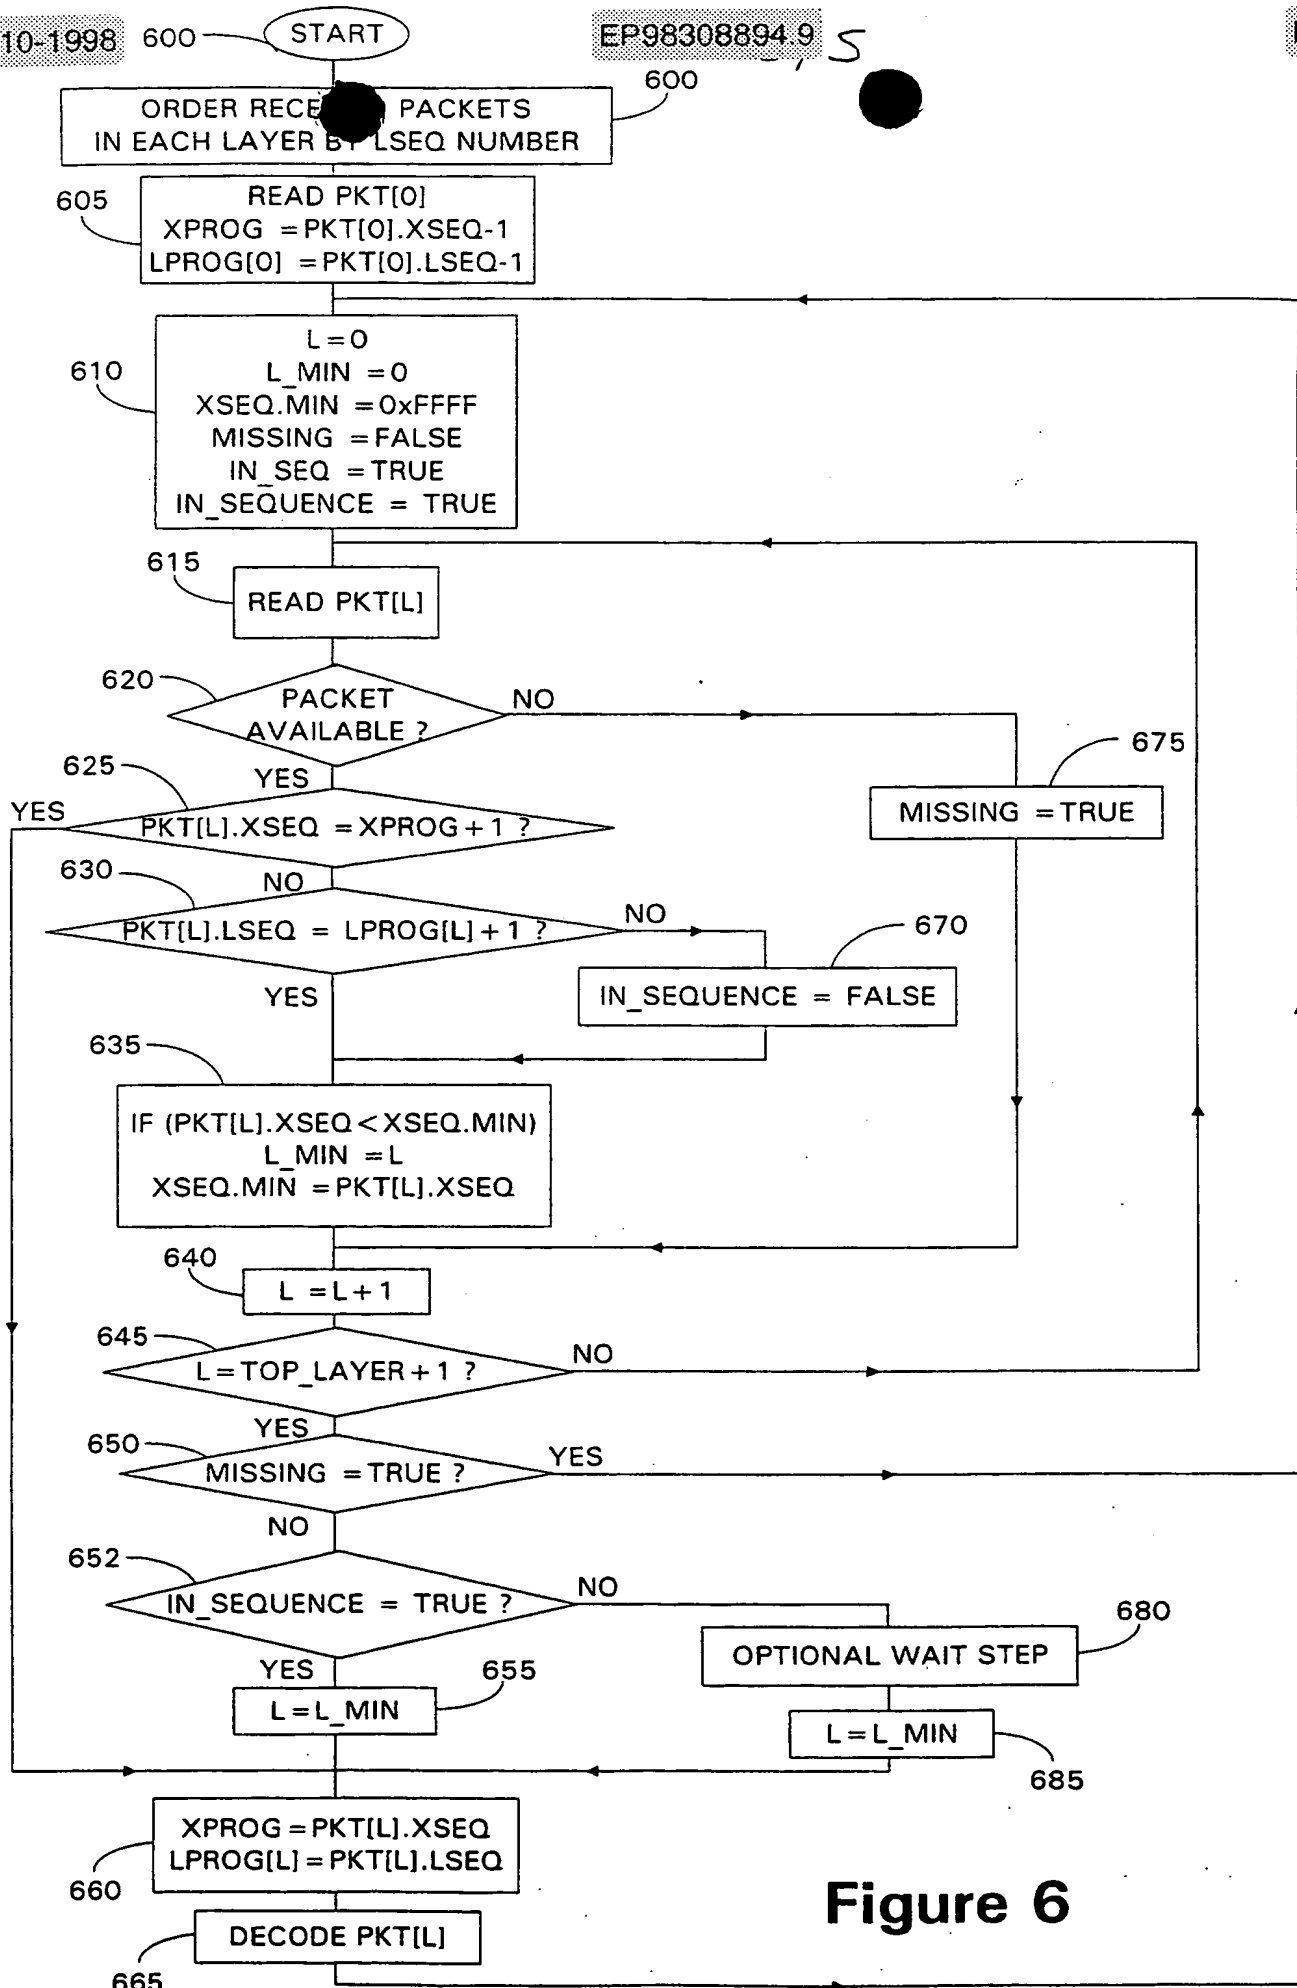


Figure 6

THIS PAGE BLANK (USPTO)